

Inferring the three-dimensional distribution of dust in the Galaxy with a non-parametric method

Preparing for Gaia

S. Rezaei Kh., C.A.L. Bailer-Jones, R.J. Hanson, and M. Fouesneau

Max Plank Institute for Astronomy (MPIA), Königstuhl 17, 69117 Heidelberg, Germany

September 29, 2016

ABSTRACT

We present a non-parametric model for inferring the three-dimensional (3D) distribution of dust density in the Milky Way. Our approach uses the extinction measured towards stars at different locations in the Galaxy at approximately known distances. Each extinction measurement is proportional to the integrated dust density along its line-of-sight. Making simple assumptions about the spatial correlation of the dust density, we can infer the most probable 3D distribution of dust across the entire observed region, including along sight lines which were not observed. This is possible because our model employs a Gaussian Process to connect all lines-of-sight. We demonstrate the capability of our model to capture detailed dust density variations using mock data as well as simulated data from the Gaia Universe Model Snapshot. We then apply our method to a sample of giant stars observed by APOGEE and Kepler to construct a 3D dust map over a small region of the Galaxy. Due to our smoothness constraint and its isotropy, we provide one of the first maps which does not show the “fingers of god” effect.

Key words. Stars: Parallaxes – Stars: Extinctions – Galaxy: Dust Map – Galaxy: ISM – Galaxy: Milky Way

1. introduction

Interstellar dust is an integral part of a galaxy. In the cycle of matter, gas and dust are ejected from evolving stars, and eventually this material will form new stars. Dust also attenuates light through reddening and extinction, thereby complicating our interpretation of stellar photometry and spectroscopy, and it imposes complex selection functions on surveys. Any survey with or of stars must account for the effects of interstellar dust.

Numerous studies have been undertaken over the years to improve our knowledge of this component of the Interstellar Medium (ISM). Early extinction maps were emission-based and in two dimensions (2D). A prominent piece of work was that of Schlegel et al. (1998) who used far-infrared dust emission measured by the IRAS and COBE satellites to build an all-sky map of the dust column density, assuming a standard reddening law. Lombardi et al. (2006) used 2MASS photometric data and the color excess technique to study molecular clouds and to map the dust column density and extinction in the Pipe nebula. Marshall et al. (2006) used a Galaxy model to find the intrinsic colors of stars, then, using the measured near infrared color excess, estimated the extinction and distances towards stars. They applied their model to more than 64 000 lines-of-sight (l.o.s) to find the 3D distribution of extinction in the inner Galaxy. Schlafly et al. (2010) used the Sloan Digital Sky Survey (SDSS) and the blue tip of the stellar locus to measure the color shift of the main sequence turnoff and thereby calculate the reddening of stars.

Sale (2012) introduced a hierarchical Bayesian model to simultaneously estimate a distance–extinction relation and the properties of individual stars from multi-band photometry. This improved the precision and accuracy of the maps compared to previous ones. This method was used by Sale et al. (2014) to

build a 3D extinction map of the northern Galactic plane from IPHAS photometry. A 3D map of interstellar dust reddening for three quarters of the sky was presented by Green et al. (2015) using Pan-STARRS1 and 2MASS photometry, based on the method of Green et al. (2014) (which uses a Bayesian approach similar to that of Sale 2012). This method was also used by Schlafly et al. (2014) to make a map of dust reddening out to 4.5 kpc from Pan-STARRS1 stellar photometry which covers the entire sky north of declination -30° . Their method was especially designed for modeling extinction in the Galactic plane.

A similar approach was taken by Hanson & Bailer-Jones (2014), who used SDSS and UKIDSS multi-band photometry to map extinction in 3D. They used a Bayesian model to take into account the degeneracy between extinction and stellar effective temperature. The method was previously introduced by Bailer-Jones (2011) to estimate effective temperatures and reddenings towards 40 000 FGK stars from 2MASS and Hipparcos photometry and Hipparcos parallaxes. Lallement et al. (2014) applied an inversion method to measurements of stellar color excess made at optical wavelengths. Together with parallaxes or photometric distances they constructed a map of the ISM within 2.5 kpc of the Sun. Sale & Magorrian (2014) presented a mapping method in which extinction is modeled as a Gaussian random field with a covariance function which has a Kolmogorov-like power spectrum. This is motivated by the idea that turbulence is responsible for the spatial structure of the ISM.

Many of the aforementioned methods consider each line-of-sight independently: they do not propagate information between neighboring lines-of-sight, even though the dust causing the extinction is likely to be correlated. This produces discontinuities in many published extinction maps, similar to an artefact known as the “fingers of god”. In this paper, we present a method which

uses extinctions and distances towards multiple stars in a collective manner simultaneously, by taking into account the correlation between neighboring lines-of-sight; unlike earlier work by Sale & Magorrian (2014), the covariance of our Gaussian process prior is in the dust density space; not in the extinction. This enables us to build a 3D map of the dust density which is free of the fingers-of-god effect, as a result of the isotropy of our smooth prior. We do this by modeling the dust density, which is a local property of the ISM, rather than the extinction, which is the integral over all the dust along the line-of-sight. We use a non-parametric model which allows us to avoid adopting an explicit – and inevitably overly simple – functional form for the variation of dust density in the Galaxy. The true variation of dust is far too complex to be captured by a parametric model, which we can both define in advance and fit well enough with available data. We instead use a Gaussian Process, which constrains the variation of the dust densities without choosing a particular functional form for its spatial variation. The Gaussian process instead defines the form of the covariance function between all points in space. This approach permits a wide range of functional variations in 3D space.

To build the dust map, we need the extinctions toward and distances of a large number of stars. Such data are being obtained by Gaia, which will soon provide astrometry and spectrophotometry for more than 10^9 stars brighter than a G-band magnitude of about 20. The expected end-of-mission parallax standard error is around $25 \mu\text{as}$ at $G = 15$ mag (de Bruijne et al. 2014), from which we can estimate distances (e.g. Bailer-Jones 2015). Gaia is also equipped with two low resolution spectrophotometers which will provide the spectral energy distribution of all observed sources. The Gaia Data Processing and Analysis Consortium (DPAC) will use these to estimate the astrophysical parameters of individual stars, including the line-of-sight reddening/extinction (Bailer-Jones et al. 2013). The end-of-mission Gaia A_V precision is expected to be around 0.03 to 0.05 mag for stars with $G \leq 15$, increasing to 0.2 mag for sources down to $G = 20$ (Andrae et al. 2016, in preparation). Individual stellar extinctions and distances are therefore the inputs for our method.

This paper is organized as follows. We introduce our method in section 2, whilst covering some technical details in the appendix. In section 3 we demonstrate the model and its ability to recover the true dust values, using both toy simulations and a simulated Gaia catalogue. We apply the model to real data over a small region of the sky (the Kepler field) in section 4. We summarize in section 5 and discuss some aspects of our method including its current strengths and weaknesses.

2. Method

We wish to determine the 3D spatial distribution of dust given measurements of the line-of-sight extinction caused by this dust toward a number of stars. Specifically, given these extinction measurements, we would like to find the probability distribution over the dust density at *any* point in space, and not necessarily a point along the line-of-sight to one of these stars.

2.1. Problem setup

Let $\rho(\mathbf{r})$ be the dust density at vector position \mathbf{r} measured from the observer. A model for the attenuation of starlight caused by this dust for a star at position \mathbf{r}_n is

$$f_n \propto \int_0^{r_n} \rho(\mathbf{r}) dr \quad (1)$$




Fig. 1. Dust cell geometry. Top: The line-of-sight towards each of the N stars is divided up into a number of dust cells. The centre of each cell is shown as a small black square. (Here we show cells of constant size.) The total number of dust cells (towards all stars) is J , and the length of dust cell j towards star n is denoted $g_{n,j}$. Bottom: These cells are represented by a sparse matrix G of size $N \times J$. Each row has non-zero elements just for the cells along the line-of-sight to that star.

where $r = |\mathbf{r}|$. The principle of our method is to invert the above to get $\rho(\mathbf{r})$ for an arbitrary point in space given measurements of the attenuation towards multiple stars. If we adopt a parametric form for $\rho(\mathbf{r})$ then this is straight forward, but the result would be highly limited by the form adopted. Here we use a non-parametric model by dividing the pencil beam along the line-of-sight toward each star into several dust cells, as shown in figure 1 (top). Let the (unknown) average dust density in cell j towards star n be $\rho_{n,j}$. The integral in equation 1 can then be replaced by a sum

$$f_n = \sum_j g_{n,j} \rho_{n,j}. \quad (2)$$

$g_{n,j}$ (the “geometric factor”) is the length of the cell along the line-of-sight, and $\rho_{n,j}$ is dust mass per unit length. Let a_n be a measurement of the attenuation towards star n . Adopting a Gaussian noise model with standard deviation σ_n means that the probability of the measurements is

$$P(a_n | \{\rho_{n,j}\}) = \frac{1}{\sqrt{2\pi}\sigma_n} \exp\left[-\frac{1}{2\sigma_n^2}(a_n - f_n)^2\right] \quad (3)$$

where $\{\rho_{n,j}\}$ denotes just those cells on the line-of-sight towards star n . Typically we measure the extinction in magnitudes, A_n , which is related to the attenuation by

$$a_n = 10^{0.4A_n} - 1. \quad (4)$$

Let us suppose that we measure the extinction toward N stars and use a total of J dust cells, for all stars. We define G as the $N \times J$ matrix with elements $g_{n,j}$, such that the n^{th} row of G contains the geometric factors just for star n . With this particular

geometry this matrix is very sparse (see figure 1), because most row elements are zero (corresponding to the cells for all other stars), and each column has just one non-zero element (stars do not share lines-of-sight). Writing the set of dust densities in all cells (for all lines-of-sight) as the J -dimensional vector ρ_J , and the model predictions for the attenuation towards the N stars as \mathbf{f}_N , we can write equation 2 as

$$\mathbf{f}_N = G\rho_J. \quad (5)$$

Writing the N attenuation measurements as the vector \mathbf{a}_N with covariance V_N , then we can generalize equation 3 to be an N -dimensional Gaussian

$$P(\mathbf{a}_N|\rho_J) = \frac{1}{(2\pi)^{N/2}|\mathbf{V}_N|^{1/2}} \exp\left[-\frac{1}{2}(\mathbf{a}_N - G\rho_J)^T \mathbf{V}_N^{-1}(\mathbf{a}_N - G\rho_J)\right]. \quad (6)$$

The above equation is the likelihood: the probability of the data given the model parameters.

Our goal is to estimate the dust density ρ_{J+1} at an arbitrary point \mathbf{r}_{J+1} in 3D space, given N measurements of the attenuation, \mathbf{a}_N , at known positions. Put probabilistically, we want to find $P(\rho_{J+1}|\mathbf{a}_N)$.

Note that each element of ρ_J refers to the average dust density in the corresponding cell, although we can consider it to be the dust density at the centre of the cell. ρ_{J+1} , in contrast, is the density at the point \mathbf{r}_{J+1} . There is no concept of a cell for points where we want to predict the dust density.

We have $J \geq N$ and $N \gg 1$. A typical problem may involve $N = 10^4$ and $J = 10^5$ (of order ten cells per star on average). In order to infer the dust densities, we need to introduce some connection between the lines-of-sight, otherwise we just have N independent equations like equation 2 with J unknowns, which would be insoluble.

2.2. Gaussian process model

We connect the lines-of-sight using a Gaussian process (e.g. Gibbs & MacKay 1997; Rasmussen & Williams 2006). This states that the joint probability distribution of the dust density at any J different points (or cell centres) is a J -dimensional Gaussian, with a covariance matrix, C_J , which depends on the distance between the points (or cell centres), i.e.

$$P(\rho_J) = \frac{1}{(2\pi)^{J/2}|C_J|^{1/2}} \exp\left[-\frac{1}{2}\rho_J^T C_J^{-1} \rho_J\right]. \quad (7)$$

We assume a zero mean Gaussian in order to have zero values for the dust density in regions where we don't have any constraints from the data. A useful property of Gaussian processes is that the conditional distribution, $P(\rho_{J+1}|\rho_J)$, is also Gaussian. A Gaussian process is just a way of specifying a prior on the covariances between points, as opposed to specifying the functional form of the dust variation in physical space (which is what a parametric model would usually do). This permits a much wider form of functional variations than a parametric model.

An important aspect of the Gaussian processes is to choose an appropriate covariance function. This determines the elements, $c_{i,j}$, of the covariance matrix between two points (or cells) i and j , with position vectors \mathbf{r}_i and \mathbf{r}_j , respectively. Here we use a covariance function from Gneiting (2002)

$$c_{i,j} = \begin{cases} \theta(1 + t^\alpha)^{-3} \left[(1 - t) \cos(\pi t) + \frac{1}{\pi} \sin(\pi t) \right] & \text{if } 0 \leq t \leq 1 \\ 0 & \text{otherwise} \end{cases}$$




Fig. 2. The covariance function in equation 8. The solid line is the case of $\alpha = 1$ (used in this work) and the dashed lines, from bottom to top, are for different values of α from 0.25, 0.5, 0.75, 1.25, 1.5, 1.75, to 2.

(8)

where

$$t = \frac{|\mathbf{r}_i - \mathbf{r}_j|}{\lambda} \quad \text{and} \quad \theta > 0, \lambda > 0.$$

which we illustrate in figure 2. We use $\alpha = 1$ (the solid line). A larger value of α (e.g. $\alpha = 2$, which has zero gradient at zero separation) produces functions with a smoother spatial variation (plots for $\alpha = 1$ are shown later). The covariance drops monotonically as the separation between the points increases, to a value of zero once the points are separated by more than λ (the scale length). Note, however, that the covariance already drops to half its maximum at $t = 0.2$, so the effective correlation distance is much less than λ . θ in the covariance function determines the overall scale of variations of the dust density. Here we consider the two hyperparameters, λ and θ , to be fixed (see section 3), although they can be inferred from the data.

This specific form of the covariance function does not have a particular physical motivation, but it has compact support, which means it goes exactly to zero beyond some distance: two points separated by more than λ will not influence each other. While other covariance functions such as an exponential will give negligible covariance beyond many scale lengths, a truncated covariance function has the advantage of giving rise to sparse covariance matrices, which reduces memory use and accelerates computations. Note that although one can easily write down any function which is truncated, this is not sufficient for it to be a covariance function (see Rasmussen & Williams 2006).

We can now use this Gaussian Process prior together with the likelihood (equation 6) to determine $P(\rho_{J+1}|\mathbf{a}_N)$. This will result in estimating $J+1$ parameters from N measurements. As $J \geq N$, this means that the resulting density estimates will not be independent. This is of course the whole point: to introduce correlations between the dust cells to make the problem tractable and – more significantly – to allow us to infer a PDF over the dust density at unobserved points.

Note that the Gaussian model allows the dust density, ρ , to be negative. Although this is unphysical, it can be tolerated and will be discussed later on.

Figure 3 shows samples drawn from the prior for different values of the hyperparameters. To make this we define a 1D grid of 1000 equally-spaced points from $r = 0$ to $r = 5000 pc$. One draw from this 1000-dimensional Gaussian gives us 1000




Fig. 3. Draws from the prior with different hyperparameters of the covariance function (equation 8). We define a 1D grid of 1000 equally-spaced points from $r = 0$ to $r = 5000pc$ and each red and black line shows one draw from a 1000-dimensional Gaussian for a fixed λ and θ .

points which are plotted at their respective positions in space, and then connected with a line. Here, for each fixed pair of λ and θ , we draw two samples from the prior (red and black lines). The sharpness of the function is due to the shape of our covariance function (fig. 2) which allows for sharper fluctuations. As mentioned earlier, a larger value of α in equation 8 produces smoother variations. It is clear that higher values of θ result in sharper and larger amplitude variations. Larger values of λ produce smoother variations. As we are using a zero-mean prior, we do see negative values for the dust density in the prior draws. This is of course unphysical, but we will see that in the presence of good data our posterior is mostly determined by the likelihood rather than the prior. We come back to this point in section 5.

2.3. Analytic solution

Using the law of marginalization over each l.o.s towards observed stars and then applying Bayes theorem, we can write the posterior PDF of the dust density at a given point given the data as

$$\begin{aligned} P(\rho_{J+1}|\mathbf{a}_N) &= \int P(\rho_{J+1}, \rho_J|\mathbf{a}_N) d\rho_J \\ &= \int \frac{P(\rho_{J+1}, \rho_J)P(\mathbf{a}_N|\rho_{J+1}, \rho_J)}{P(\mathbf{a}_N)} d\rho_J \\ &= \frac{1}{P(\mathbf{a}_N)} \int P(\rho_{J+1}, \rho_J)P(\mathbf{a}_N|\rho_J) d\rho_J, \end{aligned} \quad (9)$$

where in the last line we use the fact that \mathbf{a}_N is independent of ρ_{J+1} once conditioned on ρ_J . This is a J -dimensional integral evaluated over all values of each component of ρ_J . The term outside the integral is independent of ρ_{J+1} so is just part of the normalization constant. The first term under the integral is the Gaussian Process prior (equation 7), now in $J+1$ dimensions. The second term is the likelihood (equation 6). Both are Gaussians, but not in ρ_{J+1} . Yet because their arguments are linear functions of ρ_{J+1} , the integral will have an analytic solution. Let

$$\rho_{J+1} = \begin{bmatrix} \rho_J \\ \rho_{J+1} \end{bmatrix} \quad (10)$$

be the concatenation of the J dust densities with the dust density at the new point. Denote its covariance matrix as C_{J+1} . The distribution of ρ_{J+1} follows equation 7 with $J \rightarrow J+1$. We partition the inverse of this covariance matrix by writing it as

$$C_{J+1}^{-1} = \begin{bmatrix} \mathbf{M}_J & \mathbf{m}_J \\ \mathbf{m}_J^T & \mu \end{bmatrix} \quad (11)$$

where \mathbf{M}_J is a $J \times J$ matrix, \mathbf{m}_J is a $J \times 1$ vector, and μ is a scalar. We show in appendix A that the result of the integration is a Gaussian with mean $-\beta/\alpha$ and variance $1/\alpha$

$$P(\rho_{J+1}|\mathbf{a}_N) = \sqrt{\frac{\alpha}{2\pi}} \exp \left[-\frac{\alpha}{2} \left(\rho_{J+1} + \frac{\beta}{\alpha} \right)^2 \right], \quad (12)$$




Fig. 4. Illustration of the simple mock data in the Galactic plane. The true dust density decreases exponentially from the Galactic center (GC) in all directions. The observer is at the Sun. Region 1: in addition to the main dust variations there is a cloud between 3 and 3.5 kpc. 200 stars are observed here. Region 2: no stars are observed. Region 3: 100 stars are observed here (there is no cloud). Stars are drawn from throughout regions 1 and 3, their attenuations calculated, and observational noise added.

where

$$\begin{aligned}\alpha &= \mu - \mathbf{m}_J^T \mathbf{R}_J^{-1} \mathbf{m}_J \\ \beta &= \mathbf{a}_N^T \mathbf{V}_N^{-1} \mathbf{G} \mathbf{R}_J^{-1} \mathbf{m}_J \\ \mathbf{R}_J &= \mathbf{M}_J + \mathbf{G}^T \mathbf{V}_N^{-1} \mathbf{G}.\end{aligned}\quad (13)$$

In order to calculate the one-dimensional PDF over any single ρ_{J+1} at position \mathbf{r}_{J+1} , we must invert several matrices of size J , and this takes time $\mathcal{O}(J^n)$, where $n \lesssim 3$ for exact matrix inversion. The calculations can, however, be accelerated using certain matrix identities, as described in appendix B.

3. Demonstration using simulated data

In this section, we present some results based on both simple mock data and simulated data from Gaia Universal Model Snapshot (GUMS) (Robin et al. 2012). We use these results to explain the main properties and behaviour of the model.

3.1. Simple mock data

We first investigate the ability of our model to capture structures in the dust distribution and to infer it in unobserved regions. We further investigate the influence of the two model hyperparameters, θ and λ , on the dust density models produced.

Figure 4 illustrates the simulation set up, which specifies the distribution of dust and observations of stars in the Galactic plane. There is a general distribution of dust, the density of which decreases exponentially from the Galactic Center in all directions with a length scale of 1 kpc. Observations are made in region 1 (a 12° wedge between longitudes 354° and 6°) and region 3 (a 2° wedge between longitudes 12° and 14°), but not in region 2, a narrow wedge ($l = 7^\circ - 8^\circ$) which lies between them. In region 1 there is an additional dust cloud of 500 pc depth centered on a distance of 3.25 kpc from the Sun. This does not extend into region 3 (whether it extends into region 2 in “reality” is immaterial as we have no observations there). We want to infer the distribution of the dust density over all three regions using measurements of the attenuations towards 200 stars spread uniformly over region 1, and of 100 stars spread uniformly over region 3. Two different sets of noisy measurements are considered: standard deviations (σ_n in equation 3) of 0.1 and 0.01 on the attenuations. Figure 5 shows the noisy input data for these




Fig. 5. The attenuation in the simple mock data set as a function of distance for two different attenuation uncertainties (σ_n in equation 3) of 0.1 (left) and 0.01 (right) (no units). An increase in attenuation where the cloud is located (3 kpc to 3.5 kpc) is evident at the smaller noise level (right panel).

two situations. In region 3, the dust density increases roughly exponentially with increasing distance from the Sun. The attenuation is the integral of this, so is more or less exponential too. Region 1 is similar, but has the dust cloud in addition. This can be made out reasonably well in the data at higher signal-to-noise (right panel), but is barely noticeable at the lower signal-to-noise (left panel).

We use our model to estimate the dust density at 100 points distributed at random at each of the three regions. We set the hyperparameters to $\theta = 10^{-7}$ and $\lambda = 2$ kpc (equation 8) and use uniform cell sizes of size $g = 250$ pc. The choice of these values will be discussed later in this section. Figure 6 shows (as blue points) the estimated dust densities as a function of distance, as well as the uncertainty on this estimate (as error bars): these are the mean and standard deviation of the Gaussian posterior in equation 12. These predictions can be compared to the true values for regions 1 and 3, which are shown as red and green crosses respectively. We see that the inference of the overall exponentially-varying dust is good in all three regions, including in region 2, where there were no observations. This shows that our model performs sensible, plausible interpolations across unobserved regions. This is possible because of the smoothness prior imposed by the Gaussian prior. For region 1 (top row), the model predicts the location and density of the dust cloud well, even at the lower signal-to-noise ratio. In that case the estimated uncertainties (error bars) are also larger, which is what we want from a model. In region 2 (middle row), we ask the model to predict dust densities along a very narrow wedge between $l = 7^\circ$ and $l = 8^\circ$, located between regions 1 and 3 but closer to region 1 (with cloud) than region 3 (without cloud). With $\sigma = 0.01$ (right), an increase is obvious at 3 to 3.5 kpc corresponding to the distance of the dust cloud in region 1. A smaller increase is visible for the larger noise case (left). This too is a sensible interpolation of the available data: the cloud must stop or peter out somewhere between region 1 and region 3 because it is no longer observed in region 3. We have no information on where, but the covariance prior tells us that the closer we are to region 1, the more likely we are to still encounter the cloud. (Note that the physical transverse extent from 6° to 12° at 3 kpc is much less than the length scale, λ , we have adopted.) In region 3, the

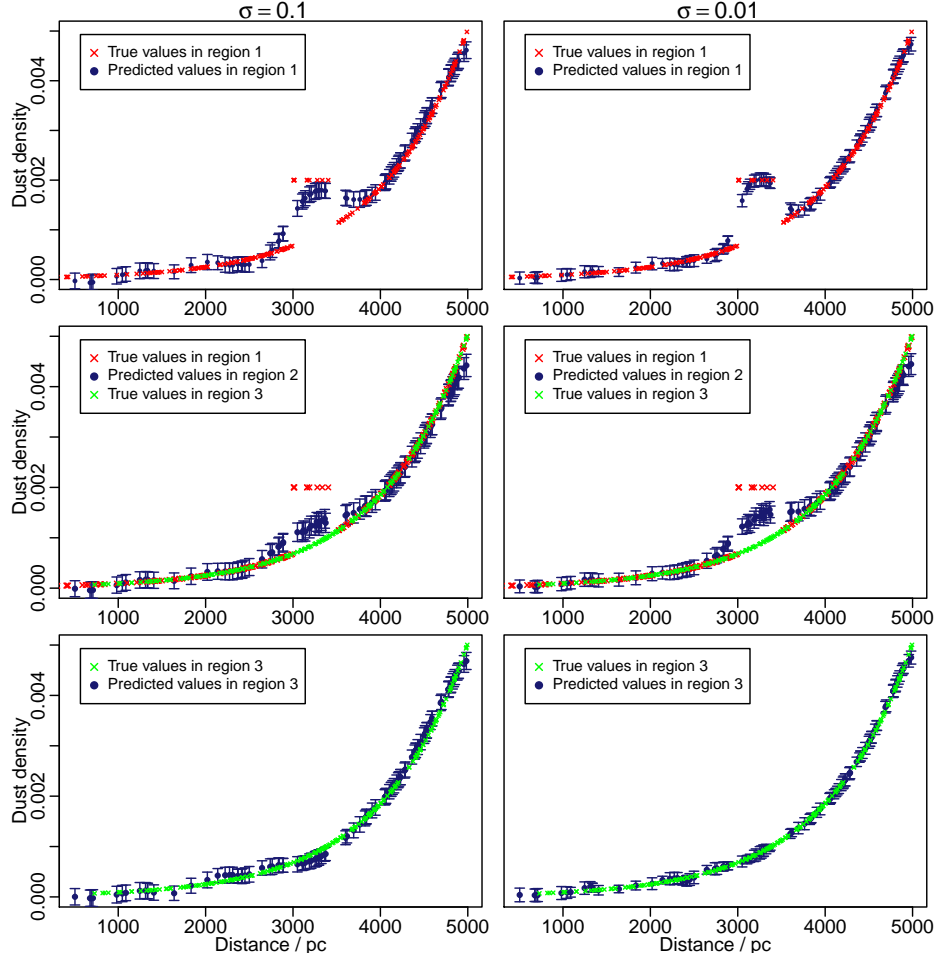


Fig. 6. Predicted dust densities in the different regions of figure 4 and for two different attenuation errors: 0.1 (left) and 0.01 (right). Red and green crosses show true values for regions 1 and 3 respectively, and blue points show predicted values for different regions (1, 2 and 3 from top to bottom). The error bars on the latter points are also predicted by our model. Although there are no observations in region 2 the model can still predict dust densities there.

model predictions show no indication of a cloud: they are influenced primarily by the nearer, cloud-free attenuation estimates.

The length scale, λ , sets the maximum distance over which dust cells are correlated. Note, however, that the correlation is only significant for values considerably smaller than λ (see figure 2). Choosing λ too small will result in too many cells being disconnected (or having very low correlations), with the outcome that the information in the data is propagated less well. λ too large will make even quite distant cells relatively highly correlated, potentially blurring out local variance determined by the data.

The cell size (here taken as constant) is the radial length over which we assume the dust density to be constant when setting up the model. It is only used by equation 2 (generally equation 5) to discretize the dust density for representing the dust attenuation towards observed stars. It is used neither in the calculation of the covariance nor in the computation of dust density at new points, so contrary to possible expectations it does not represent the minimum scale over which we can compute density variations. It does, however, set some kind of minimum length scale over which we are sensitive to dust variations. Ideally we would use very small cells, but the computation time grows as the third power of the number of cells, so in practice we are limited by computational considerations.

θ sets the scale of the covariance and thus the amplitude of variations in the dust. For a given λ and cell size, a larger value of θ means we can capture larger variations in the dust. We see from equation 8 with $t = 0$ that θ is the expected variance in the dust at any point. An estimate for the value of θ is therefore the variance in the expected distribution of the dust density over all cells. We can get an order of magnitude estimate of this before applying the model, by using the simplifying assumption that towards a given star, n , every cell has the same dust density, ρ_n , and same variance therein, $\text{Var}(\rho_n)$. Adopting a constant cell size g_n for a given star, it follows from equation 2, using a_n as our estimate of f , that

$$\rho_n = \frac{a_n}{j_n g} \quad (14)$$

where $j_n g$ is just the distance to the star. Let μ_ρ be the average of the $\{\rho_n\}$ across all N stars. The weighted variance of this distribution is

$$\theta = \frac{\sum_{n=1}^N w_n (\rho_n - \mu_\rho)^2}{\sum_{n=1}^N w_n} \quad (15)$$

where each weight, w_n , can be set equal to the inverse of the variance, $\text{Var}(\rho_n)$, in the corresponding value of ρ_n . These can be

found by taking variance equation 2 (with a_n as our estimate of f) in which $\rho_{n,j} = \rho_n$ is consistent with our above assumptions. This gives

$$\text{Var}(\rho_n) = \frac{\text{Var}(a_n)}{j_n g_n^2} \quad (16)$$

where j_n is the number of cells towards the star n . Equations 14–16 allow us to estimate θ using the measured attenuations, a_n , and their uncertainties, $\sigma = \sqrt{\text{Var}(a_n)}$ as well as the distances and adopted cell sizes. We will see that the larger the value of θ , the larger the error bars on our dust density predictions will be, as we expect. Using the simple mock data, we get $\theta = 1.9 \times 10^{-7}$ for $\sigma = 0.1$ and $\theta = 1.8 \times 10^{-7}$ for $\sigma = 0.01$.

It's important to realize that the minimum length scale over which we can probe dust variations is *not* set by λ . We see in figure 7 and we will see later that we can probe significant variations on length scales much less than λ . The finite cells sizes (g) aside, the minimum length variation is actually set by a combination of λ and θ , plus, most importantly, by the data themselves.

Figure 7 shows the effects of varying the values of the hyperparameters on the model predictions. Here we compare the true and estimated dust density for different values of λ (columns) and θ (rows) for region 1 with $\sigma = 0.1$. As before, red crosses show true values and blue points show the predictions. λ and θ increase from left to right and top to bottom, respectively. Larger values of λ produce smoother variations in the dust. This is largely because this connects more cells, thereby increasing the amount of data used to estimate the dust densities. This also decreases the uncertainties on the estimates (the error bars). Smaller values of θ prevent the model from following steep changes in the dust density: for given λ we see smoother variations in the dust density at smaller θ .

The very nature of our covariance model is that points separated by less than λ have correlated posterior distributions. Thus not only are the predicted dust estimates correlated, but so are their predicted uncertainties. Thus the error bars of adjacent points in figure 7 are highly correlated. They are simple point estimates of a continuous function (we could make estimate on a grid twice as dense; the error bars would not change). Significant variations of the dust density can be (and are) obtained which are far smaller than these error bars.

In a real application – where we don't know the true dust densities – we could compute the likelihood of the data at the predicted values of the dust densities from equation 6. We could do this for a range of hyperparameters λ and θ and find the model value which gives the highest likelihood. This is a form of Bayesian model selection, because the regularizing prior has been used to estimate the individual dust densities. The different models correspond to different values of the hyperparameters.¹ The one practical disadvantage of this is that we would have to compute the dust densities along the line-of-sight to every star and for all values of hyperparameters, which could become very time consuming.

¹ What we call the likelihood here is not the likelihood in the sense of parametric models, which is the probability of the data for a given model *and* a given set of parameters. While we could maximize this likelihood to find the best parameters for a given model, we cannot use it to choose among models, because it contains no regularization and so will just identify the most complex model (we can eventually fit the data perfectly, noise and all). The likelihood we are talking about in this Gaussian Process context is at a higher level. The “parameters”, if you will, have effectively been marginalized over by the Gaussian Process to produce the probability of the data for given hyperparameters.

For the sake of this simulation – where we do know the true dust densities – we compute instead the probability of the inferred dust densities given the true dust densities. By construction, the inferred dust densities have a joint Gaussian distribution, and the true dust densities have no variance. The log probability is therefore given by

$$\ln P = -\frac{1}{2} \Delta \boldsymbol{\rho}^T \mathbf{C} \Delta \boldsymbol{\rho} - \frac{1}{2} \ln((2\pi)^n |\mathbf{C}|), \quad (17)$$

which we will call the *reconstruction probability*. It is analogous to the likelihood (or negative log of the sum-of-squared residuals, if the covariance were unity) but with model predicted values replaced by their true ones. n is the number of predicted points, \mathbf{C} is the $n \times n$ covariance matrix of these points with elements given by equation 8, and $\Delta \boldsymbol{\rho}$ is the $n \times 1$ vector of the differences between true and predicted dust densities. The numbers in the corners of the panels in figure 7 show the values of this metric. The highest value is at $\lambda = 2000$ pc and $\theta = 1 \times 10^{-7}$, which are the values we used for our predictions in figure 6. Recall that our pre-modeling order of magnitude estimate of what to use for θ gave a value of $\theta = 1.9 \times 10^{-7}$ for $\sigma = 0.1$, which is very similar. This procedure couldn't specify λ , because we are free to specify this according to the flexibility of the fitting we wish to achieve (and/or our knowledge of the true scale of the variations). For a fixed θ of 10^{-7} and different values of λ , we see from figure 7 that as long as λ is large enough (a few times cell sizes) to connect many cells in 3D space ($\lambda = 1000, 2000$ and 3000 pc in this example), we can achieve good results.

From figure 7, it is clear that the model predicts the dust densities with smaller uncertainties when using larger values of λ . This is because it then uses more points to predict the dust densities for every new point, although the closer points of course still have more influence (correlation) than more distant points. As the dependence on λ is not strong (once θ is set), we choose to fix λ to a relatively large value (a few times the cell size).

3.2. Gaia Universe Model Snapshot (GUMS)

We now look at a more realistic set of simulations taken from the Gaia Universe Model Snapshot (GUMS) (Robin et al. 2012).

GUMS is a simulation, generated by the DPAC prior to the Gaia launch, of what the Gaia catalogue can be expected to contain. It contains both the intrinsic properties of the objects as well as simulations of the noise-free Gaia observations (or more precisely, the corresponding catalogue products). It comprises around 1.6×10^9 stars (in single or multiple systems) with G-band magnitudes brighter than 20. We select the coordinates, distance, G-band magnitude, absolute V-band magnitude, $(V - I)$ color, extinction A_V , and effective temperature T_{eff} for these stars. GUMS uses a dust model to generate its extinction values, but this is not part of our simulated catalogue (we have no knowledge of the true dust densities).

To map dust extinction we need not – and should not – use all stars. Due to computational limitations, the model, as it stands, cannot cope with anything nearly as large as the number of objects in the catalogue (see section 5 for possible improvements). We therefore first select just those stars with $G < 15$, parallax errors less than 5%, and T_{eff} between 5000 and 10000 K. Using cooler stars, e.g. down to 3000 K, we found that it does not make a significant difference in the results, but as their parallax and/or extinction estimates would often be less precise, we would probably omit them in practice. The magnitude selection is imposed to ensure that the Gaia spectrophotometry have high SNR so that A_V is determined to better than about 0.05 mag by

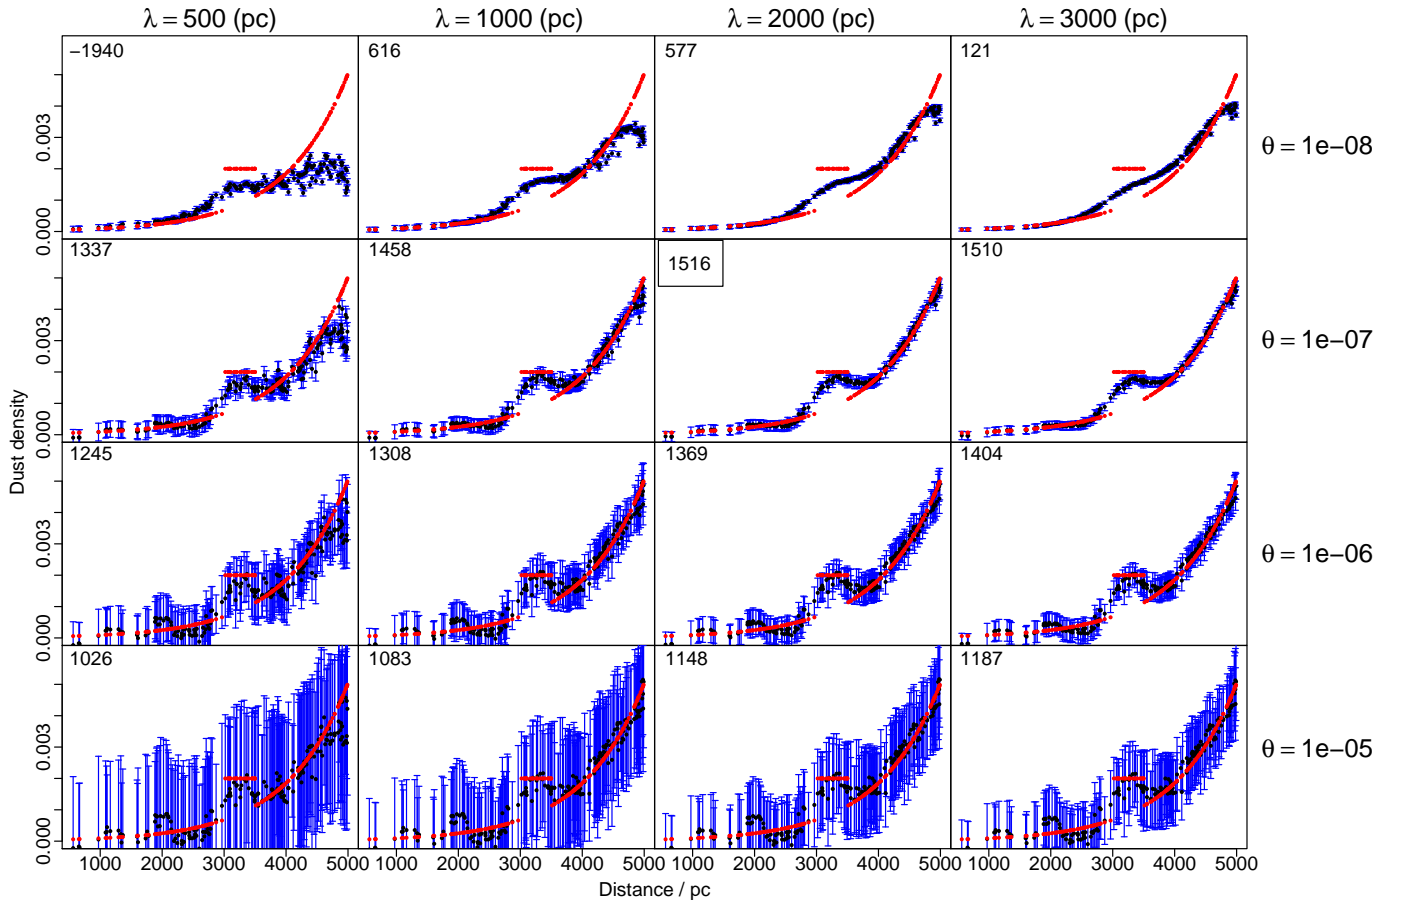


Fig. 7. The effects of the hyperparameters on the model predictions in region 1 with $\sigma = 0.1$. Red points are the true values and black points with blue error bars are predictions by the model. Note that the predicted uncertainties are highly correlated. Each column is for different values of λ (length scale), and each row is for different values of θ . The number in the top left corner is the natural logarithm of the reconstruction probability, defined by 17.

the DPAC processing (Bailer-Jones et al. 2013) and we calculate the expected end-of-mission parallax errors from the simplified formula in de Bruijne et al. (2014). Although a selection on apparent magnitude biases our sample towards less extinct stars, this is hard to avoid in practice (as almost all surveys have a magnitude limit). As Gaia provides parallaxes from which we can infer distances, we could instead attempt to select all stars within a given volume. But this selection (magnitude and temperature limit) is better because the distance uncertainties are asymmetric and are themselves a strong function of magnitude, plus the interstellar extinction would significantly limit the size of a complete volume at low Galactic latitudes. A better approach might be to limit the selection to intrinsically bright stars, over a narrow T_{eff} range (and therefore absolute magnitude range) for which we can estimate accurate extinctions. This will be explored in more detail in subsequent work.

Applying the above selections gives us 30 million stars distributed throughout the simulated Galaxy. To ease the interpretation of the results for the sake of this demonstration, we select stars in a wedge narrow in longitude (5° to 7°) but broad in latitude (-1° to 30°) within 2 kpc containing around 52 000 stars. Their positions and extinctions are shown in figure 8. The drop off at high latitudes is due to the decline in density of the Galactic disk population away from the plane. The presence of small, nearby regions of higher dust density are apparent from the higher attenuations along some lines-of-sight.




Fig. 8. Two-dimensional positions of the 52 000 stars meeting our selection criteria over the region $5^\circ < l < 7^\circ$ and $-1^\circ < b < 30^\circ$ out to 2 kpc from the GUMS catalogue, color-coded by the true (GUMS) extinctions in magnitudes.

A useful feature of our method is that we do not have to use all stars within a region. Although the stars are independent




Fig. 9. A random set of 1000 stars selected from those shown in figure 8. See section 3.2.




Fig. 10. Predictions of the dust densities for the selected area from GUMS catalogue ($5^\circ < l < 7^\circ$ and $-1^\circ < b < 30^\circ$) for 2000 new points in two dimension colored coded by the dust density values. A high value dust cloud is seen at the large distances and low latitudes as well as some lower-value dust clouds in other locations.

probes of their line-of-sight extinctions, stars in close proximity to one another probe much of the same dust. Provided the stars retained have a high enough spatial density to map out the minimum scale of the spatial variations we want to probe, randomly removing additional stars will not qualitatively change the resulting map. It may reduce the precision, because we then have fewer measurements to determine the dust density, but often other factors dominate the uncertainties (we will examine this in section 5). However, as reducing the number of stars can lead to great computational savings and increased numerical stability, this is a useful strategy to pursue. Figure 9 shows the positions of 1000 stars randomly selected from the sample of 52 000. Although there are far fewer stars, we can still see most of the structures from the full sample.

Figure 10 shows the predicted dust densities for 2000 random new points in the selected region. A localized region of high dust density (“dust cloud”) is apparent at around $(X, Z) = (400, 150)$ which is responsible for the high extinctions beyond this point visible in figure 9, namely the upper diagonal wedge. Likewise, the very high extinction region in the plane beyond about 1400 pc is assigned a cloud at the same distance. Less apparent is a dense dust cloud responsible for the intermediate diagonal wedge in figure 9; the model has instead attributed this to a more diffuse region of higher dust density, perhaps because there is a larger region of higher extinctions at intermediate latitudes. Note that the position of these new points have been chosen at random; we are free to select them and predict dust density at any point in 3D space. For this predictions we used a dust correlation length scale, λ , of 2 kpc, dust cell lengths, g , of 250 pc and a dust variation scale, θ , of 4×10^{-7} .

In the above we selected just 1000 stars from a possible 52 000 which met our selection criteria. How might this selection affect our inference? The top panel of Figure 11 shows the same results as in Figure 10 but now giving the average dust density in the rectangular region. The middle panel shows the results when we instead select 1000 stars but biased towards selecting higher attenuations. For the biased selection, we sample from the data with the probability of being selected proportional to the rank of the sorted attenuation values. The bottom panel is for 2000 stars selected in this same biased manner. Overall we see a high level of consistency, although not surprisingly, these latter two data sets reveal larger densities in the higher extinction

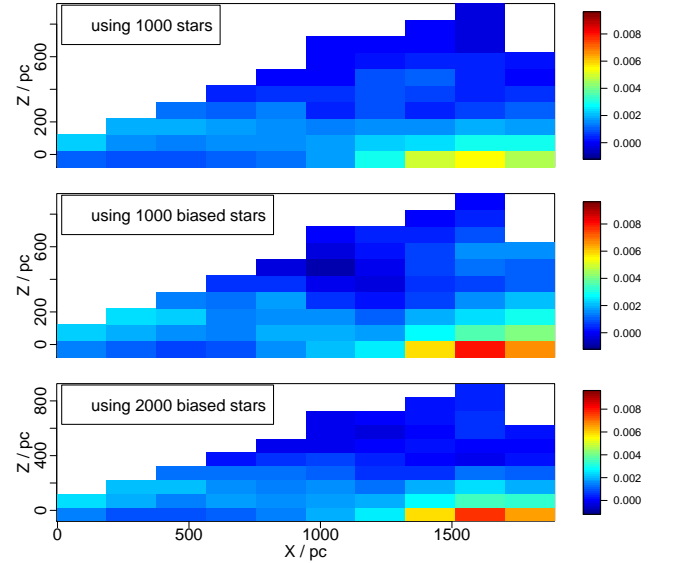


Fig. 11. The effect of different sampling on the predictions for the region of $5^\circ < l < 7^\circ$ and $-1^\circ < b < 30^\circ$ using GUMS catalogue (section 3.2). Top panel shows the predictions using 1000 randomly sampled input data (from the entire 52 000), the middle panel is the results using 1000 input stars which sampled in a biased way towards high value attenuations, and the lower panel shows predictions in case of the same biased sampling but using 2000 stars. The results are pretty much consistent, capturing similar trends.

areas. There is essentially no difference between using 1000 and 2000 stars, however. This is good news for our method, because it has a poor computational time scaling with the number of data points.




Fig. 12. Extinction values of 1900 stars from Rodrigues et al. (2014) lying within 3 kpc. The color indicates the extinction in magnitudes. The discrete squares are due to the Kepler satellite fields.




Fig. 13. Inferred dust density for the area of $68^\circ < l < 85^\circ$ and $6^\circ < b < 20^\circ$ out to 3 kpc using data from Rodrigues et al. (2014) shown in 2D with the same frame as figure 12 (the input data). Note that dust density is local in 3D and these points lie at a range of distances as seen in figure 15.

4. Application to APOKASC data

Having demonstrated the basic features of our model on simulated data, we now apply it to a set of real data to construct a 3D dust map in a small region of the Galaxy.

We use extinctions and positions of nearly 2000 stars provided by Rodrigues et al. (2014). They use spectroscopic and asteroseismic data of giants observed by APOGEE and the Kepler satellite (APOKASC catalogue), together with photometry from SDSS, 2MASS, and WISE and apply a Bayesian method to determine their extinctions and distances. The sample covers 17 degrees in longitude (from 68° to 85°) and 14 degrees in latitude (from 6° to 20°). We select stars with distance uncertainties of less than 10% within 3 kpc and remove stars with negative extinctions. This leaves around 1900 stars with observed extinctions shown in figure 12. The spatial distribution arises from the Kepler’s CCD placement in the focal plane. Many of the extinctions are small, but we do see patches of higher extinction on different scales, especially in the lower right of the plot.

We apply our model to estimate the dust density (specifically: the mean and standard deviation of a Gaussian distribution at each point) at 2000 points distributed at random throughout the volume occupied by the stars. This has the shape of a pyramid with its apex at the Sun. We use $\lambda = 2$ kpc, $g = 250$ pc, and $\theta = 4 \times 10^{-7}$ (set as described in section 3.1). Figure 13 shows the mean estimated dust density of each of these points projected onto the sky. Comparing the coverage of this dust map with the input data makes it apparent that we can predict dust densities for points outside of the measured field. The model predicts an extended region of higher density at the bottom of the region, presumably driven by the higher extinctions for some stars in the lower fields in figure 12.

As the inferred dust distribution covers a range of distances, a sky projection like Figure 13 does not give the full picture. Figure 14 attempts to show the three-dimensional distribution. As expected from the input data, regions with higher dust densities are located at the lower latitudes, as is visible at the bottom of the east and west triangles. There is also a higher density region at about 1.5 kpc from the Sun which, as it is most visible in the south triangle, must be located predominantly in the southern part of the region. Note that there is no line-of-sight (“fin-

gers of god”) effect in the dust reconstruction. This is a direct consequence of the non-parametric nature of our model plus its minimal assumption on a smoothness prior. This is one of the advantages that our approach brings over more traditional mapping techniques, which are generally based on independent estimates of the dust density along different lines-of-sight, which lead necessarily to the fingers of god.

One way to assess the performance of our model is to estimate the dust density at several points along the line-of-sight towards a star, and then to use equation 2 to predict the attenuation, f_n . The standard deviation in this prediction is found by taking the variance of equation 2,

$$\text{Var}(f_n) = \mathbf{g}^T \mathbf{C} \mathbf{g} \quad (18)$$

where \mathbf{g} is the the vector of cell sizes along that line-of-sight, and \mathbf{C} is the covariance matrix (with elements given by equation 8) of the dust densities in these cells. Note that this expression takes into account the (often large) covariance between cells along a line-of-sight. (This is essential, because our model by its very nature assumes spatial correlations in the dust.) We compute attenuation estimates and uncertainties in this way for 200 stars in the two lower right Kepler fields in figure 12. For each star we use 15 cells with constant cells sizes per line-of-sight. As the stars are at a side range of distances, this corresponds to cells sizes between 39 and 334 pc.

Figure 15 shows the residuals (predicted minus measured) of the attenuations to these 200 stars, along with the corresponding error bars. These error bars are computed from the sum of the variance in the measurement, σ_d^2 , and the variance from equation 18. We see no particular trend in the magnitude of either the residuals or their uncertainties with distance. However, it looks as though there is a small negative bias, in the sense that our model slightly under estimates the measured attenuations. This is better seen in Figure 16, which shows the residuals scaled by their uncertainty estimates. A linear model with ideal Gaussian residuals would show a Gaussian distribution with zero mean and unit standard deviation. A negative bias is apparent. This is not necessarily a sign of a bad model, because the whole point of the model is to make an inference from the data subject to the

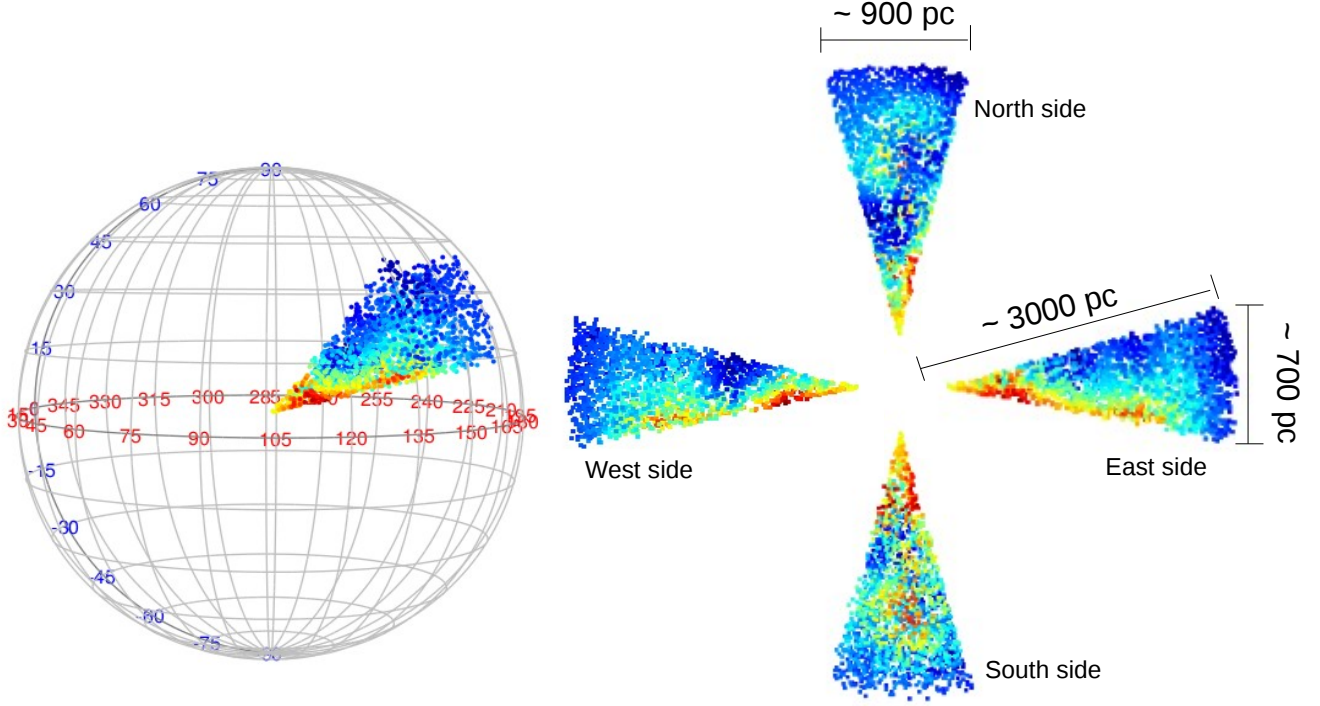


Fig. 14. Dust densities estimated by our model over the pyramid-shaped region $68^\circ < l < 85^\circ$, $6^\circ < b < 20^\circ$, $d < 3$ kpc, computed using the extinction and distance data from Rodrigues et al. (2014). The left panel figure shows the data in three dimensions. The right panel shows the dust density in projection as viewed from the outside sides of the pyramid. 'North' means looking from top to bottom, perpendicular to the upper side of the pyramid, and similarly for 'south' (looking from the bottom). The 'East' and 'West' views are looking perpendicular to the sides of the pyramid, where 'east' means looking from higher longitudes to lower ones, and 'west' from lower to higher. Note that these triangles are plotted with a larger scale parallel to their base (transverse to the line-of-sight from the Sun) – a larger opening angle than in reality – in order to better resolve the details.

smoothness constraint. Our model uses a zero mean prior for the dust density; thus, in the absence of data the model gives zero density. We can always add an offset to the covariance function of the Gaussian process to have a non-zero mean (e.g. if we want to examine the high density regions).

In contrast, the standard deviation is almost exactly unity, indicating that our Gaussian Process is very good at estimating the uncertainty in its predictions. Figure 17 shows the predicted values together with the measured values as a function of distance. Recall that the error bars on the predictions are correlated. They increase with distance because the more distant stars have larger dust cells. There are a few measurements which lie well outside of the main envelope of the data, but given the number of points (and the size of their error bars in some cases) this is entirely consistent. We see how the model has smoothed out these "outliers", on account of its built-in assumptions. We further see that our model's underestimation of high extinctions is a consequence of its smooth variation with distance.

5. Summary and future improvements

We have introduced a new non-parametric method for building a smooth, three-dimensional map of dust opacity which avoids line-of-sight effects. It uses a Gaussian Process prior to constrain the variation of the dust density in 3D space, but without assuming a specific functional form for the spatial dependence. It instead uses a covariance function which varies with the separation of points. This allows the model to infer the dust density in unobserved regions. Our model uses the 3D positions of stars together with their line-of-sight extinctions as its input data and infers the posterior probability density function (PDF) over the dust at selected points. This PDF is a Gaussian, and we showed that its mean and standard deviation have analytic solutions. While the lines-of-sights to the observed stars are divided into discrete cells, predictions are made at arbitrary points without any discretization being necessary.

We used a truncated covariance function in the Gaussian process which involves two hyperparameters: a correlation length scale λ and a dust amplitude θ . The latter can be set from the




Fig. 15. Residuals between reconstructed and measured attenuations as a function of distance for the 200 stars in the APOKASC data (from the bottom right two fields shown in Figure 12).




Fig. 16. Distribution of the scaled residuals, $(a_n - f_n)/\sigma$, where $\sigma^2 = \sigma_a^2 + \text{Var}(f_n)$, for the set of 200 stars in the APOKASC data. The mean is -0.21 and the standard deviation is 0.99. The curve shows a Gaussian with these values for comparison.

properties of the input data; adjusting λ gives us flexibility to model dust variations on different length and amplitude scales. The only requirement is that λ be larger than the cell sizes. While λ is some characteristic length scale, our model can and does probe dust structures of much smaller scales (as in figure 7).

We could try to fix these hyperparameters by calculating the Bayesian evidence. The evidence (or “marginal likelihood”) is the probability of observing the data, for fixed λ and θ , averaged over all possible instantiations of the model. We compute this by drawing one sample from the J -dimensional Gaussian process prior (which gives us J values for the dust density), calculating the likelihood for these model dust densities, repeating it for a large number of times (e.g. $K = 10^5$), and then averaging these likelihoods;

$$P(\mathbf{a}_N | \lambda, \theta) = \frac{1}{K} \sum_{k=1}^K P_k(\mathbf{a}_N | \{f_i\}, V_N), \quad (19)$$

where $\{f_i\}$ are calculated attenuations (equation 5) using dust densities drawn from the prior. Having done this for various λ and θ , we then calculate the Bayes factors, which are the ratio of these evidences (for different λ and θ) to the one with λ and θ used for our APOKASC data (section 4; $\lambda = 2$ kpc and $\theta = 4 \times 10^{-7}$). We report these in table 1.

We get values for the Bayes factors in the case of APOKASC data (section 4) which agree broadly with what we calculated for λ and θ . But in the case of the simulated data (section 3.1), the Bayesian evidence does not give us a useful discrimination between models. Most values are very close to zero for a range of θ and λ because our simulated data have high extinctions, which are not well represented by a Gaussian process prior with zero mean. The APOKASC data, in contrast, have smaller extinctions. This shows that using a non-zero mean in the Gaussian process prior will better construct the dust density in regions

Table 1. Logarithm of the Bayes factors for APOKASC data for different ranges of λ and θ .

$\text{Log}_{10}(\text{Bayes factor})$				
$\theta \setminus \lambda$ (pc)	500 pc	1000 pc	2000 pc	3000 pc
1×10^{-8}	-3.30	-2.64	-1.83	-1.42
1×10^{-7}	-0.76	-0.28	-0.03	-0.01
1×10^{-6}	+0.04	-0.01	-0.10	-0.18
1×10^{-5}	-0.29	-0.39	-0.57	-0.62

with higher extinctions, such as the disk of the galaxy and the spiral arms. One could also use a different covariance function in the Gaussian process. We tested various forms of the covariance function, such as truncated exponential forms, but they did not make a significant difference to our results. The covariance function that we are using has the advantage that we can get different variation slopes by changing α (see fig. 2).

One drawback of our model is that its computation time increases nonlinearly with the number of stars, N , and the number of cells, J . As explained in appendix B, the time-consuming part is the (one-off) inversion of the $J \times J$ covariance matrix C_J , which takes time $O(J^n)$ to compute, where n is typically $\lesssim 3$ but can be reduced to around 2.3 (Demmel et al. 2007), as well as various matrix inversions and multiplications taking time $O(NJ^2)$, which must be done for every prediction. For a problem with $N = 230$ and $J = 3203$, inverting C_J took two minutes (using a single core on a modest AMD Opteron 6380 CPU). Making predictions at multiple points can then be done in parallel: 200 predictions took 4 minutes with 40 cores, or 1.2 seconds per point. The computation time for more points is proportional to the number of points. For a problem with $N = 1000$ and $J = 8185$, it took 40 minutes to invert C_J , and 34 seconds per point to make new predictions for 1000 new points (again with 40 cores). This




Fig. 17. Reconstructed attenuations as a function of distance (black dots with blue error bars) on top of the measured values (in red) for the set of 200 stars in the APOKASC data. Note that the error bars in the predictions are highly correlated.

is 28 times longer than the previous case, which agrees reasonably well with the $O(NJ^2)$ scaling suggested above (which gives $(1000 \times 8185^2)/(200 \times 3203^2) = 33$).

The limiting factor when scaling this up to larger applications may be the memory rather than the run-time. For the case of $J = 12\,000$, we needed 8 GB of RAM per core. This number is determined primarily by the number of cells, J , because the largest matrix has size $J \times J$. However, as we use sparse matrix methods and a truncated covariance function, the RAM required will not continue to grow as J^2 . It is rather the density of cells in space, rather than the number of cells, which will ultimately drive the memory requirements.

Using the run-time numbers from above, and ignoring memory limitations, then with $N=10\,000$ and $J=100\,000$, the C_J inversion takes around 30 days (with just one core; this could be accelerated if C_J inversion is parallelized too). Then even with 10 000 cores running for 30 days we could only make predictions at 30 000 points. This (and N) is too small to build up a useful dust density map over a large volume of space. One way to accelerate the computations is to use approximate matrix inversions, which can be done in time $O(N^2)$. Alternatively, instead of trying to model the entire volume in one, we could partition it into partially overlapping regions, solve for each separately, and then join them. Thought is required to combine the overlapping regions without discontinuities, but recall that any two points separated by more than λ are not connected by our model anyway, due to the truncated covariance matrix (which itself does not produce discontinuities). Optimal partitioning is an area for future investigation.

While our method takes into account the uncertainties in the extinction measurements, it does not yet make use of the distance uncertainties. This will be necessary in some practical applications, as even from Gaia more distant and/or fainter stars will

have poor distance estimates from the parallaxes (e.g. Bailer-Jones 2015). Including distances in the likelihood model is one approach, but would make the solution non-analytic. We are currently exploring this approach and will report on it in a future paper.

Once these developments have been made and tested, the model will be ready to be applied to the Gaia data.

Acknowledgments

We wish to thank David Hogg, Rene Andrae, and Hans-Walter Rix for their useful discussions and suggestions. This project is partially funded by the Sonderforschungsbereich SFB 881 “The Milky Way System” of the German Research Foundation (DFG).

References

- Bailer-Jones, C. A. L. 2011, MNRAS, 411, 435
- Bailer-Jones, C. A. L. 2015, PASP, 127, 994
- Bailer-Jones, C. A. L., Andrae, R., Arcay, B., et al. 2013, A&A, 559, A74
- de Bruijne, J. H. J., Rygl, K. L. J., & Antoja, T. 2014, in EAS Publications Series, Vol. 67, EAS Publications Series, 23–29
- Demmel, J., Dumitriu, I., & Holtz, O. 2007 [arXiv:0612264]
- Gibbs, M. & MacKay, D. J. 1997
- Gneiting, T. 2002, Journal of Multivariate Analysis, 83, 493
- Green, G. M., Schlafly, E. F., Finkbeiner, D. P., et al. 2014, ApJ, 783, 114
- Green, G. M., Schlafly, E. F., Finkbeiner, D. P., et al. 2015, ApJ, 810, 25
- Hanson, R. J. & Bailer-Jones, C. A. L. 2014, MNRAS, 438, 2938
- Lallement, R., Vergely, J.-L., Valette, B., et al. 2014, A&A, 561, A91
- Lombardi, M., Alves, J., & Lada, C. J. 2006, A&A, 454, 781
- Marshall, D. J., Robin, A. C., Reylé, C., Schultheis, M., & Picaud, S. 2006, A&A, 453, 635
- Rasmussen, C. E. & Williams, C. 2006
- Robin, A. C., Luri, X., Reylé, C., et al. 2012, A&A, 543, A100
- Rodrigues, T. S., Girardi, L., Miglio, A., et al. 2014, MNRAS, 445, 2758
- Sale, S. E. 2012, MNRAS, 427, 2119
- Sale, S. E., Drew, J. E., Barentsen, G., et al. 2014, MNRAS, 443, 2907
- Sale, S. E. & Magorrian, J. 2014, MNRAS, 445, 256
- Schlafly, E. F., Finkbeiner, D. P., Schlegel, D. J., et al. 2010, ApJ, 725, 1175
- Schlafly, E. F., Green, G., Finkbeiner, D. P., et al. 2014, ApJ, 789, 15
- Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, ApJ, 500, 525

Appendix A: Analytic solution of the integral

To solve the integral in equation 9, we write it as

$$P(\rho_{J+1}|\mathbf{a}_N) = \frac{1}{Z} \int e^{-\psi/2} d\rho_J, \quad (\text{A.1})$$

where Z is a normalization constant, and where from equation 7 (with $J \rightarrow J+1$) and equation 6

$$\begin{aligned} \psi &= \rho_{J+1}^\top C_{J+1}^{-1} \rho_{J+1} + (\mathbf{a}_N - G\rho_J)^\top V_N^{-1} (\mathbf{a}_N - G\rho_J) \\ &= \rho_{J+1}^\top C_{J+1}^{-1} \rho_{J+1} + \mathbf{a}_N^\top V_N^{-1} \mathbf{a}_N - \rho_J^\top G^\top V_N^{-1} \mathbf{a}_N + \\ &\quad \rho_J^\top G^\top V_N^{-1} G \rho_J - \mathbf{a}_N^\top V_N^{-1} G \rho_J \end{aligned} \quad (\text{A.2})$$

with

$$\rho_{J+1} = \begin{bmatrix} \rho_J \\ \rho_{J+1} \end{bmatrix}. \quad (\text{A.3})$$

As each term in equation A.2 is a scalar, and transposing a scalar leaves it unchanged, we see that the third and fifth terms are identical.

To evaluate the integral we need to separate ρ_J from ρ_{J+1} in ρ_{J+1} . We first partition the inverse covariance matrix in the following way

$$C_{J+1}^{-1} = \begin{bmatrix} \mathbf{M}_J & \mathbf{m}_J \\ \mathbf{m}_J^\top & \mu \end{bmatrix} \quad (\text{A.4})$$

where \mathbf{M}_J is a $J \times J$ matrix, \mathbf{m}_J is a $J \times 1$ vector, and μ is a scalar. Note that these are components of the inverted matrix, not components of C_{J+1} which are then inverted. We can then write

$$\begin{aligned} \rho_{J+1}^\top C_{J+1}^{-1} \rho_{J+1} &= \begin{bmatrix} \rho_J^\top & \rho_{J+1} \end{bmatrix} \begin{bmatrix} \mathbf{M}_J & \mathbf{m}_J \\ \mathbf{m}_J^\top & \mu \end{bmatrix} \begin{bmatrix} \rho_J \\ \rho_{J+1} \end{bmatrix} \\ &= \rho_J^\top \mathbf{M}_J \rho_J + 2\rho_{J+1} \mathbf{m}_J^\top \rho_J + \mu \rho_{J+1}^2. \end{aligned} \quad (\text{A.5})$$

Substituting this into equation A.2 and gathering together terms gives

$$\begin{aligned} \psi &= \rho_J^\top (\mathbf{M}_J + G^\top V_N^{-1} G) \rho_J + 2(\rho_{J+1} \mathbf{m}_J^\top - \mathbf{a}_N^\top V_N^{-1} G) \rho_J \\ &\quad + (\mathbf{a}_N^\top V_N^{-1} \mathbf{a}_N + \mu \rho_{J+1}^2), \end{aligned} \quad (\text{A.6})$$

which is a quadratic expression in ρ_J . The last term is independent of ρ_J so can be taken out of the integral, allowing us to write equation A.1 as

$$\begin{aligned} P(\rho_{J+1}|\mathbf{a}_N) &= \frac{1}{Z} \exp \left[-\frac{1}{2} \mathbf{a}_N^\top V_N^{-1} \mathbf{a}_N \right] \exp \left[-\frac{1}{2} \mu \rho_{J+1}^2 \right] \\ &\quad \int \exp \left[-\frac{1}{2} \rho_J^\top \mathbf{R}_J \rho_J + \mathbf{b}_J^\top \rho_J \right] d\rho_J, \end{aligned} \quad (\text{A.7})$$

where

$$\mathbf{R}_J = \mathbf{M}_J + G^\top V_N^{-1} G \quad (\text{A.8})$$

$$\mathbf{b}_J = G^\top V_N^{-1} \mathbf{a}_N - \rho_{J+1} \mathbf{m}_J. \quad (\text{A.9})$$

The integral is a standard one

$$\begin{aligned} \int \exp \left[-\frac{1}{2} \rho_J^\top \mathbf{R}_J \rho_J + \mathbf{b}_J^\top \rho_J \right] d\rho_J \\ = \frac{(2\pi)^{J/2}}{|\mathbf{R}_J|^{1/2}} \exp \left[\frac{1}{2} \mathbf{b}_J^\top \mathbf{R}_J^{-1} \mathbf{b}_J \right] \end{aligned} \quad (\text{A.10})$$

provided $|\mathbf{R}_J| > 0$. Using this we can write equation A.7 as

$$P(\rho_{J+1}|\mathbf{a}_N) = \frac{1}{Z} e^{-\phi/2} \quad \text{where} \quad \phi = \mu \rho_{J+1}^2 - \mathbf{b}_J^\top \mathbf{R}_J^{-1} \mathbf{b}_J \quad (\text{A.11})$$

and we have absorbed all factors which do not depend on ρ_{J+1} into the normalization constant. Substituting for \mathbf{b}_J from equation A.9 this becomes

$$\begin{aligned} \phi &= -\mathbf{a}_N^\top V_N^{-1} G \mathbf{R}_J^{-1} G^\top V_N^{-1} \mathbf{a}_N + (\mu - \mathbf{m}_J^\top \mathbf{R}_J^{-1} \mathbf{m}_J) \rho_{J+1}^2 + \\ &\quad 2\mathbf{a}_N^\top V_N^{-1} G \mathbf{R}_J^{-1} \mathbf{m}_J \rho_{J+1}. \end{aligned} \quad (\text{A.12})$$

The first term does not depend on ρ_{J+1} so can also be absorbed into the normalization constant. Putting this into equation A.11 gives us

$$P(\rho_{J+1}|\mathbf{a}_N) = \frac{1}{Z} \exp \left[-\frac{1}{2} \alpha \rho_{J+1}^2 - \beta \rho_{J+1} \right], \quad (\text{A.13})$$

where

$$\begin{aligned} \alpha &= \mu - \mathbf{m}_J^\top \mathbf{R}_J^{-1} \mathbf{m}_J \\ \beta &= \mathbf{a}_N^\top V_N^{-1} G \mathbf{R}_J^{-1} \mathbf{m}_J. \end{aligned} \quad (\text{A.14})$$

By completing the square in the exponent

$$-\frac{1}{2} \alpha \rho_{J+1}^2 - \beta \rho_{J+1} = -\frac{\alpha}{2} \left(\rho_{J+1} + \frac{\beta}{\alpha} \right)^2 + \frac{\beta^2}{2\alpha} \quad (\text{A.15})$$

and taking the term $\exp(\frac{\beta^2}{2\alpha})$ outside of the integral (and absorbing it too into normalization constant), we see that $P(\rho_{J+1}|\mathbf{a}_N)$ is just a Gaussian with mean $-\beta/\alpha$ and variance $1/\alpha$, i.e.

$$P(\rho_{J+1}|\mathbf{a}_N) = \sqrt{\frac{\alpha}{2\pi}} \exp \left[-\frac{\alpha}{2} \left(\rho_{J+1} + \frac{\beta}{\alpha} \right)^2 \right]. \quad (\text{A.16})$$

Appendix B: Accelerating the matrix evaluations via matrix identities

We can simplify the expressions for α and β and thereby reduce the number of matrix multiplications using some standard matrix identities. We first relate the components of the inverted covariance matrix (equation A.4) to the components of the covariance matrix itself, which we partition as

$$C_{J+1} = \begin{bmatrix} C_J & \mathbf{k}_J \\ \mathbf{k}_J^\top & k \end{bmatrix}. \quad (\text{B.1})$$

C_J is the $J \times J$ covariance matrix involving only the fixed data. \mathbf{k}_J is the $J \times 1$ vector with elements given by $c_{j,J+1}$, i.e. the covariance between the fixed data and the new point. k is the scalar $c_{J+1,J+1}$, the covariance of the new point with itself (i.e. its variance). Recall that each element of the covariance matrix is determined by the covariance function, such as equation 8. A standard result of matrix algebra (e.g. Press et al. 1992) allows us to write the components of the inverted matrix as

$$\mathbf{M}_J = C_J^{-1} + (C_J^{-1} \mathbf{k}_J)(k - \mathbf{k}_J^\top C_J^{-1} \mathbf{k}_J)^{-1} (\mathbf{k}_J^\top C_J^{-1}) \quad (\text{B.2})$$

$$\mathbf{m}_J = -(C_J^{-1} \mathbf{k}_J)(k - \mathbf{k}_J^\top C_J^{-1} \mathbf{k}_J)^{-1} \quad (\text{B.3})$$

$$\mu = (k - \mathbf{k}_J^\top C_J^{-1} \mathbf{k}_J)^{-1}. \quad (\text{B.4})$$

Temporarily defining

$$\mathbf{h}_J = C_J^{-1} \mathbf{k}_J \quad (\text{B.5})$$

allows us to write the matrix and vector parts of the inverted covariance matrix as

$$\mathbf{M}_J = \mathbf{C}_J^{-1} + \mu \mathbf{h}_J \mathbf{h}_J^\top \quad (\text{B.6})$$

$$\mathbf{m}_J = -\mu \mathbf{h}_J \quad (\text{B.7})$$

Note that $\mathbf{h}_J \mathbf{h}_J^\top$ is an outer product, so produces a matrix of size $J \times J$. We now use the matrix inversion lemma (also known as the Woodbury matrix identity; e.g. Press et al. 1992) to write, from equation A.8,

$$\mathbf{R}_J^{-1} = \mathbf{M}_J^{-1} - \mathbf{M}_J^{-1} \mathbf{G}^\top (\mathbf{V}_N + \mathbf{G} \mathbf{M}_J^{-1} \mathbf{G}^\top)^{-1} \mathbf{G} \mathbf{M}_J^{-1}. \quad (\text{B.8})$$

The Sherman–Morrison formula (a special case of the matrix inversion lemma) allows us to write

$$\begin{aligned} \mathbf{M}_J^{-1} &= \mathbf{C}_J - (1 + \mu \mathbf{h}_J^\top \mathbf{C}_J \mathbf{h}_J)^{-1} \mu \mathbf{C}_J \mathbf{h}_J \mathbf{h}_J^\top \mathbf{C}_J \\ &= \mathbf{C}_J - (1 + \mu \mathbf{k}_J^\top \mathbf{C}_J^{-1} \mathbf{k}_J)^{-1} \mu \mathbf{k}_J \mathbf{k}_J^\top \end{aligned} \quad (\text{B.9})$$

where we have substituted for \mathbf{h}_J from equation B.5.

All of the above allows us to simplify our expressions for α and β in equation A.14 by using

$$\begin{aligned} \mathbf{R}_J^{-1} &= \mathbf{M}_J^{-1} - \mathbf{M}_J^{-1} \mathbf{G}^\top (\mathbf{V}_N + \mathbf{G} \mathbf{M}_J^{-1} \mathbf{G}^\top)^{-1} \mathbf{G} \mathbf{M}_J^{-1} \\ \mathbf{M}_J^{-1} &= \mathbf{C}_J - (1 + \mu \mathbf{k}_J^\top \mathbf{C}_J^{-1} \mathbf{k}_J)^{-1} \mu \mathbf{k}_J \mathbf{k}_J^\top \\ \mathbf{m}_J &= -\mu \mathbf{C}_J^{-1} \mathbf{k}_J \\ \mu &= (k - \mathbf{k}_J^\top \mathbf{C}_J^{-1} \mathbf{k}_J)^{-1} \\ k &= c_{J+1, J+1} \\ \mathbf{k}_J &= (c_{1, J+1}, c_{2, J+1}, \dots, c_{J, J+1})^\top. \end{aligned} \quad (\text{B.10})$$

\mathbf{C}_J^{-1} takes time $O(J^n)$ to compute, where n is typically $\lesssim 3$, but must only be done once. As we use a truncated covariance function, \mathbf{C}_J is a sparse matrix, a fact which we exploit when inverting it. To compute the PDF at each new point, various other matrix inversions and multiplications must be performed, the longest of which takes time $O(NJ^2)$, because $J \geq N$. Note that \mathbf{G} is always sparse (when $J = N$ it is even diagonal, although this would give very poor resolution dust density maps). Because \mathbf{C}_J is sparse, \mathbf{M}_J^{-1} is also sparse.